REMARKS

This Amendment is responsive to the Final Office Action dated August 8, 2008. Applicant has amended claims 1, 4, 9, 11, 12, 13, 14 and 16. Applicant has cancelled claims 7 and 18. Applicant has added claims 26-32. Claims 1-6, 8-17, and 19-32 are pending upon entry of this amendment.

Amendment to the Specification

Applicant has amended the specification to correct a typographical error. Before entry of this amendment, Applicant's paragraph [0027] in the published application stated:

> FIG. 4 illustrates one exemplary software architecture for implementing the principles of the invention. More specifically, the software architecture includes a number of cooperating software modules. The embodiments of the exemplary software architecture may reside on each of the N servers originating a data transfer to a client device. In particular, portions or all of the software modules of FIG. 4 may be loaded on each source server.

However, Applicant's claim 1 as presented at the time of filing through to this current amendment specifies that that invention is executed on a client computer. For example, the preamble of claim 1 states, "a parallel download system executing on a client computer." In contrast to previously presented paragraph [0027], the parallel download system is not executed on the servers, but is instead executed on the client device.

Applicant has amended paragraph [0027] to remove any indication that the parallel download system, e.g., the exemplary software module of FIG. 4, resides on the server(s). No new matter is added by this amendment.

Claim Rejection Under 35 U.S.C. § 112

In the Final Office Action, the Examiner rejected claims 1, 3, 4, 7, 8, 14, and 16 under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. Particularly, the Examiner asserted that the Applicant failed to point out where the limitations of the amended claims are supported in the specification, citing to MPEP chapter 2163.03 sec. I, chapter 2163.04 sec. I, and chapter 2163.06.

Applicant respectfully traverses the 35 U.S.C. § 112, first paragraph rejections. The following cites to various paragraphs within Applicant's published application that provide explicit support for Applicant's claim amendments presented in the response to the previous Office Action.

Claim 1

Applicant has substantially amended claim 1. As amended, the § 112, first paragraph rejection is moot. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 1.

Claim 3

Applicant refers the Examiner to paragraph [0076] of Applicant's published application. Paragraph [0076] specifies that the prioritization scheduler comprises a data prioritizer, and that the data prioritizer determines the priority of the data to be scheduled from a highest priority to a lowest priority. Paragraph [0076] in combination with paragraph [0075] provides explicit support for the phrase "wherein different portions of the file are downloaded in parallel from the two or more of the plurality of servers in accordance with the ordering of the data within the file as specified by the data prioritzer and the download schedule as dynamically adjusted by the prioritization scheduler," which was added by amendment in the response to the previous Office Action.

Accordingly, no new matter was added by the amendment to claim 3. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 3.

Claim 4

Applicant refers the Examiner to paragraph [0074] and [0032] of Applicant's published application. As noted above, in paragraph [0074] the prioritization scheduler tracks the latency and throughput of all channels at all times. Paragraph [0032] states that the invention allows any number of channels to be opened to any combination of sources. Accordingly, in one non-limiting example, all channels are coupled to combinations of sources, i.e. servers. Paragraphs [0074] and [0032] provide explicit support for the phrase "prioritization scheduler adjusts the

download schedule with respect to the different servers while maintaining the prioritization order in which the data within the file will be received," as required by claim 4.

Accordingly, no new matter was added by the amendment to claim 4. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 4.

Claims 7 and 8

Applicant has cancelled claim 7 in the current response. Accordingly, the 35 U.S.C. § 112, first paragraph rejection for claim 7 is moot.

Claim 8 was amended to correct minor grammatical errors. For example, before amendment, claim 8 recited a bulk scheduler to determines which bytes of data. Applicant amended claim 8 to delete the extraneous "s" at the end of determines, since the proper grammatical form should be determine.

Accordingly, no new matter was added by the amendment to claim 8. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 8.

Claim 14

Applicant refers the Examiner to paragraph [0123] of Applicant's published application. Applicant had amended claim 14 in response to the Examiner's recommendation in the previous Office Action. Particularly, Applicant deleted the phrase "such as a Merkle Hash tree." Paragraph [0123] specifies that a simple integrity verifier uses a full file hash. Paragraph [0123] further specifies that the full file hash algorithms can be calculated incrementally. Paragraph [0123] provides explicit support to the phrase "the integrity verification engine utilizes a hash construct," as required by claim 14.

Accordingly, no new matter was added by the amendment to claim 14. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 14.

Claim 16

Applicant is confused by the 35 U.S.C. § 112, first paragraph rejection with respect to claim 16. Applicant did not amend claim 16 in the response to the previous Office Action. Nevertheless, Applicant refers the Examiner to paragraph [0114] of Applicant's published

application. Paragraph [0114] specifies that it is the channel connector's job to establish a new channel to the scheduled source. Paragraph [0114] provides explicit support to the phrase "one or more channel connector to establish a new channel to the scheduled source server," as required by claim 16.

Applicant did not amend claim 16, accordingly, no new matter was added. Applicant respectfully requests withdrawal of 35 U.S.C. § 112, first paragraph rejection for claim 16.

Claim Rejection Under 35 U.S.C. § 103

In the Final Office Action, the Examiner rejected claims 1–13, 15–18 and 20–25 under 35 U.S.C. § 103(a) as being unpatentable over Young (US 6,477,522) in view of Byers et al. ("Accessing Multiple Mirror Sites in Parallel", April 1999) and rejected claims 14 and 19 under 35 U.S.C. § 103(a) as being unpatentable over Young, in view of Byers, and further in view of Merkle (US 4,881,264). Applicant respectfully traverses the rejection to the extent the rejection may be applicable to the amended claims. The applied references fail to disclose or suggest the inventions defined by Applicant's claims, and provide no rational reason that would have suggested the desirability of modification to arrive at the claimed invention.

Claim 1

Applicant has amended claim 1 to further clarify the invention. The applied references fail to teach or suggest various features of claim 1 as amended. For example, the applied references fail to teach or suggest "a source scheduler that ranks a plurality of communication channels to generate a download schedule to control which of the plurality of communication channels will be downloaded from in parallel at a point in time, wherein each of the source servers stores a copy of at least a portion of a file containing data, and wherein each of the plurality of communication channels comprises a network connection to one of the plurality of source servers," as required by claim 1.

Textual support for this feature of claim 1 may be found in Applicant's published application at paragraphs [0032] and [0058]. Paragraph [0058] clarifies that the source scheduler ranks the channel sources, e.g., communication channels. Moreover, paragraph [0032] in Applicant's published application clarifies that the invention allows any number of channels to

be opened to any combination of sources. Paragraph [0032] provides textual support that the feature that each of the plurality of communication channels comprises a network connection to one of the plurality of servers.

As another example, the applied references fail to teach or suggest "a prioritization scheduler that determines a range of the data that should be requested from one or more of the plurality of source servers and the point in time when the range of data should be requested from one or more of the plurality of source servers, wherein the prioritization scheduler tracks an overall latency and throughput for all the communication channels for one or more of the source servers while downloading at least a portion of the data of the file from two or more of the plurality of communication channels in parallel and, based on the latency and throughput of the communication channels, dynamically adjusts the download schedule including the range of data and the point in time to request the range of data for at least one of the communication channels while downloading the data in parallel to control download performance," as now required by claim 1.

Textual support for this feature may be found at paragraphs [0074] and [0081] of Applicant's published application. Paragraph [0074] states that the prioritization scheduler will decide what data is requested from the sources and at what time. Paragraph [0081] states the prioritization scheduler schedules ranges of bytes, e.g., range of data, to be downloaded. Paragraph [0074] also explicitly states that the prioritization scheduler tracks latency and throughput of all channels to adjust the download schedule.

As yet another example, the applied references fail to teach or suggest "a proportional allocator that, for a source server of the one or more source servers, determines a plurality of portions of the range of data for that source server that should be downloaded using one or more of the communication channels to that source server, wherein each of the portions of the range of data to be downloaded from that source server is determined based on one or more of an expected throughput and latency of each of the one or more communication channels to that source server and a time interval assigned by the proportional allocator for downloading the range of data using the communication channels," as now required by claim 1.

Textual support for this feature may be found at paragraph [0090] of Applicant's published application. Paragraph [0090] states that the proportional allocator determines how

much data will be transferred from the channel that is being scheduled. Paragraph [0090] makes clear that the proportional allocator determines how much of the data of the range of data set by the prioritization scheduler for one or more of the source servers is downloaded from each of the communication channels, e.g., the channel being scheduled. Stated another way, the proportional allocator will determine how much of the range of data should be downloaded from each of the channels. The prioritization scheduler has already determined the range of data that should be downloaded from the one or more source servers. The proportional allocator provides further granularity and determines how much data of the range of data should be downloaded from each one of the one or more servers.

Paragraph [0090] also states that this amount of data will be proportional to the throughput that is expected to be received over the channel. Paragraph [0090] also states that the proportional allocator will use time intervals to determine the proportion of data to be downloaded. Paragraph [0090] makes clear that the proportional allocator determines the portion of the range of data that is downloaded by each of the communication channels based on throughput and a time interval.

Young

Claim 1 requires that each of the plurality of communication channels comprises a network connection to one of the plurality of source servers. Contrary to the requirements of claim 1, Young does not disclose a plurality of communication channels that provide a network connection to a plurality of source servers. Instead in Young, the computing device only establishes one connection to one server at any given time. The Examiner recognized this deficiency in Young.

Claim 1 also requires that prioritization scheduler, based on the latency and throughput of the communication channels, dynamically adjusts the download schedule including the range of data and the point in time to request the range of data for at least one of the communication channels while downloading the data in parallel to control download performance.

Contrary to the requirements of claim 1, in addition to not teaching parallel downloads, Young also does not adjust the download schedule including the range of data and the point in time to request the range of data. In Young, the range of the data and the time when the range of

data is downloaded is fixed. In Young, the computing device first downloads a fixed range of data, e.g., 500 bytes, from the first server. The computing device then downloads the next 500 bytes from the second server, and so on. The range of the data is completely fixed and is never adjusted. Furthermore, the point in time when the range of data is requested to be downloaded is also fixed. In Young, the download of data is sequential, i.e., the first fixed range of data is downloaded, and then the next (second) fixed range of data is downloaded, and so on. Accordingly the point in time when the range of data is downloaded is fixed. It will always follow the time when the previous range of data was downloaded.

Contrary to Young, in claim 1, the prioritization scheduler dynamically adjusts the range of data and the point in time when the range of data is requested. In Young, the range of data is fixed while in claim 1, the range of data is dynamic and can be adjusted during the actual download of the data. Furthermore, in Young not only can the range of data not be adjusted, but it cannot be adjusted during the download as required by claim 1. For example, in Young when the computing device is downloading from the first server, it cannot dynamically adjust the range of data it is receiving from the first server during download from the first server. In this respect, and contrary to claim 1, not only is the range of data fixed during the download from any server, the range of data is fixed for all the servers until the computing device determines the best server. However, claim 1 requires that the prioritization scheduler should be able to dynamically adjust the range of data, and further more, the prioritization scheduler should be able to dynamically adjust the range of data during the download.

Claim 1 also requires that a proportional allocator, for a source server of the one or more source servers, determine a plurality of portions of the range of data for that source server that should be downloaded using one or more of the communication channels to that source server.

Contrary to Young, the proportional allocator provides finer granularity to the range of data downloaded from each server. In accordance with claim 1, the prioritization scheduler determines the range of data that should be downloaded from one or more of the servers. The proportional allocator determines the portion of the range of data that should be downloaded from each server for each communication channel for that server.

Young does not contemplate a proportional allocator. Young make no reference to dividing the range of data into a plurality of portions of data. In Young there would be no need

to include a plurality of portions of the range of data because Young provides only one connection to a server at a given time. Accordingly, in Young there would be no necessity nor any benefit to include a proportional allocator.

For at least these reasons, Young fails to teach or suggest the features of claim 1.

Byers

As noted above, claim 1 requires that the prioritization scheduler, based on the latency and throughput of the communication channels, dynamically adjusts the download schedule including the range of data and the point in time to request the range of data for at least one of the communication channels while downloading the data in parallel to control download performance.

Contrary to the requirements of Byers, claim 1 requires that the prioritization scheduler adjust the range of data and the point in time to request the range of data for at least one of the communication channels. Byers discloses a system where the parallel mirror sites establish multicast groups to transmit data.[1] In Byers, all the data is transmitted simultaneously from all mirror sites to the client device. When the client device receives *k* number packets from the servers, the client device is capable of reconstituting the entire data. In this respect and contrary to claim 1, Byers does not contemplate a range of data that is transmitted via a communication channel or a point in time to request the range of data. Instead, in Byers, all the data is transmitted at the same time until *k* number of packets have been downloaded by the client device. Accordingly, and in contrast to claim 1, Byers would have no need for a prioritization scheduler because all the data in Byers is transmitted simultaneously.

As noted above, claim 1 also requires that a proportional allocator, for a source server of the one or more source servers, determine a plurality of portions of the range of data for that source server that should be downloaded using one or more of the communication channels to that source server.

In contrast to claim 1, Byers does not disclose a proportional allocator, nor would there be a need for a proportional allocator. As described above, Byers multicasts all of the data simultaneously to the client device. As noted above, Byers does not contemplate a range of data,

[1] Byers, first page.

and accordingly, cannot teach or suggest a plurality of portions of the range of data, as required by claim 1.

For at least these reasons, Byers fails to teach or suggest the features of claim 1.

As described above, both Young and Byers fail to teach or suggest virtually all the features of amended claim 1. Accordingly, Young in view of Byers fails to teach or suggest each and every feature of claim 1. Applicant respectfully submits that claim 1 patentably distinguishes Young in view of Byers. For at least the reasons advanced above, Applicant respectfully requests withdrawal of the 35 U.S.C. § 103(a) rejection for claim 1.

Furthermore, Applicant submits that one of ordinary skill in the art contemplating Young and Byers would not have modified Young based on the teachings of Byers. Instead, as explained in more detail below, one of ordinary skill in the art contemplating Young and Byers would have rejected the teaching of Young, and instead only implemented the techniques of Byers.

Young teaches a system where a computing device determines which server provides the best throughput based on downloading different portions of a file from different servers, or downloading the same portion of a file from different servers.[2] Young makes clear that the Young system sequentially downloads a portion of the file from different servers[3]. In a further embodiment, the computing device checks the throughput from the selected server.[4] If the throughput from the selected server drops below a threshold, the computing device either selects the already determined next best server, or recalculates for the next best server, and continues to download only from the next best server.[5]

Byers explicitly rejects the techniques of Young as being inadequate and inefficient. Byers teaches many problems associated with techniques disclosed in Young. Applicant refers the Examiner to the section entitled "Inadequacy of a Basic Scheme," on the third page of the Byers reference. In the "Inadequacy of a Basic Scheme" section, Byers describes a basic scheme that is equivalent to the scheme taught by Young. Particularly, Byers states to divide the file into

[2] Young, col. 2, lines 5-15.
[3] Young, Summary.
[4] Young, col. 2, lines 16-18.
[5] Young, col. 2, lines 18-22.

equally sized disjoint blocks.[6] And then download a first block from one of the servers, after the sender finishes sending a block, the renegotiation stage would request a block that has not be sent.[7] This teaching of Byers is substantially similar to Young's technique of sequentially downloading a portion of the file from different servers.

Byers criticizes Young's techniques of adjusting which server to download from based on the throughput rate. Byers specifically states that if transmission rates, i.e. throughput rate, vary significantly overtime, several renegotiations may be required to obtain the file at the fastest rate.[8] Byers asserts that these problems grow more significant with the number of senders a receiver uses, suggesting that this approach, while implementable, may not scale well to many senders.[9] In essence, Byers teaches away from Young's techniques and instead teaches a different methodology of receiving data.

As described above, Byers fundamentally rejects the techniques of Young as being inefficient and inadequate. Importantly, Byers fundamentally rejects sequentially downloading blocks of data from separate servers, and fundamentally rejects dynamically determining which server provides the best throughput rate, i.e. transmission rate, and downloading from that server. Accordingly, one of ordinary skill in the art contemplating the techniques and teachings of Young and Byers would reject the techniques of Young, because of the teachings of Byers, and only implement the teachings of Byers.

As previously described, one of ordinary skill in the art contemplating the techniques and teachings of Young and Byers would reject the techniques of Young because of the teachings of Byers. Accordingly, modifying Young in view of Byers means rejecting the techniques of Young and only implementing the teachings of Byers. As also described above, Byers fails to teach or suggest various features of claim 1. For at least these additional reasons, Applicant respectfully requests withdrawal of the 35 U.S.C. § 103(a) rejection for claim 1 based on Young in view of Byers.

Dependent Claims

[6] Byers, third page.
[7] Byers, third page.
[8] Byers, third page.
[9] Byers, third page.

The claims dependent on independent claim 1, namely claims 2-6, 8-17, and 19-25, incorporate all of the limitations of claim 1, and therefore are patentable for at least the reasons expressed above. In light of the shortcomings of Young and Byers with respect to independent claim 1, Applicant reserves comment with respect to the dependent claims. In reserving comment, Applicant does not acquiesce to the Examiner's interpretations or assertions regarding the dependent claims. Applicant respectfully requests withdrawal of the 35 U.S.C. § 103(a) rejection for dependent claims 2-6, 8-13, 15-17, and 20-25.

Claim 3

Furthermore, Applicant had previously argued that claim 3 requires that the data prioritizer specifies an ordering of the data within the file from a highest priority data to a lowest priority data, and wherein different portions of the file are downloaded in parallel from the two or more of the plurality of servers in accordance with the ordering of the data within the file as specified by the data prioritzer and the download schedule as dynamically adjusted by the prioritization scheduler. Young does not provide teaching of a system that prioritizes and schedules the download of different portions of a file in parallel from different servers.

The Examiner found Applicant's arguments with respect to claim 3 to be non-persuasive. Particularly, the Examiner asserted that "lowest priority" is interpreted as being the initial file portions (for examples, blocks of 500 bits) that are downloaded from the servers. The Examiner asserted that "highest priority" would be the remaining data within the file that is downloaded from the server with highest throughput, performance, etc. The Examiner cited to Young col. 4, lines 39-67 and col. 5, lines 13-25 in support of the argument.[10]

Young does not specify a priority for the blocks of data within a file. Instead, Young merely sequentially downloads the blocks of data. Keeping the Examiner's example, and the example provided in the Young, the client device downloads the first 500 bits, e.g. 1-500, from a first server, then downloads the next 500 bits, e.g. bits 501-1000, from a second server, and repeats the process for all the servers. Finally, based on the throughput of each of server, the client devices selects the server with the highest throughput rate and sequentially downloads the remaining bits of the data file.

[10] Final Office Action, page 3.

In contrast to Young, claim 3 requires "wherein different portions of the file are downloaded in parallel from the two or more of the plurality of servers in accordance with the ordering of the data within the file as specified by the data prioritzer and the download schedule as dynamically adjusted by the prioritization scheduler." According to the requirements of claim 3, the data prioritizer specifies the order in which to download the various blocks within the file. As one non-limiting example, keeping with the Examiner's example, in accordance with claim 3, the data prioritizer may determine that bits 1-500 need to be downloaded from the first server, and bits 1001-1500, need to downloaded in parallel from the second server. Bits 1-500 and 1001-1500 may be downloaded first in parallel, i.e. bits 1-500 and 1001-1500 are given the highest priority. The data prioritizer may also determine that bits 501-1000 needs to be downloaded from the fourth server and bits 1501-2000 need to be downloaded from the fifth server. Bits 501-1000 and bits 1501-2000 may be downloaded only after the bits 1-500 and 1001-1500 are downloaded, i.e. bits 501-1000 and bits 1501-2000 are given the lowest priority. Accordingly, as claimed in claim 3, different portions of the file are downloaded in parallel from the two or more of the plurality of servers in accordance with the ordering of the data within the file as specified by the data prioritzer and the download schedule as dynamically adjusted by the prioritization scheduler. In contrast, to claim 3, Young makes no determination of the highest priority and the lowest priority. Instead, Young merely downloads sequentially. In Young, and contrary to claim 3, there would be no need for a data priortizer since the data is downloaded sequentially.

Stated simply, Young does not specify a priority of the data blocks. The data blocks are merely downloaded in a sequential fashion. Accordingly, in accordance with Young, there would be no need for a data prioritizer that specifies a priority of the data blocks. In contrast to Young, and as required by claim 3, the data prioritizer assigns the download priority to each of the blocks of data such that the client device may not download the blocks of data in a sequential fashion.

Claim 11

Claim 11 requires a constraint scheduler to ensure that the source scheduler does not attempt to retrieve data that a source server cannot provide. In support of the rejection, the

Examiner cited to Young col. 4 line 58 to col. 5 line 9. Young at col. 4 line 58 to col. 5 line 9 describes switching which server the client device is downloading from when the throughput rate falls below some threshold. Contrary to claim 11, Young at col. 4 line 58 to col. 5 line 9 makes absolutely no mention of ensuring that the source scheduler does not attempt to retrieve data that a source server cannot provide. Nowhere in Young is there any teaching or suggestion of ensuring a source server can provide the data. Young assumes that each of the source servers can provide the data.

Claim 21

Claim 21 requires the data prioritizer to adjust the ordering of the data within the file based on which one or more external applications are attempting to access the data within the file. In contrast to claim 21, Young is completely silent regarding adjusting the ordering of the data within the file because as described above Young merely downloads the data sequentially. In Young there would be no need to adjust the ordering of the data. Accordingly, Young fails to teach or suggest a data prioritizer that adjusts the ordering, much less, a data prioritizer that adjusts the ordering of the data based on which one or more external applications are attempting to access the data within the file. Young is wholly silent on external applications, much less, external applications that are used by the data prioritizer to determine the ordering of the data within the file.

Claims 14 and 19

Moreover, in the Final Office Action, the Examiner rejected claims 14 and 19 under 35 U.S.C. § 103(a) as being unpatentable over Young in view of Merkle (US 4,881,264). Applicant respectfully traverses the rejection. Merkle fails to provide any additional teachings to overcome the deficiencies of Young and Byers. Claim 14 and 19 are allowable insofar as they are dependent upon allowable base claim 1. Applicant reserves the right for further comment regarding claims 14 and 19. Applicant respectfully requests withdrawal of the 35 U.S.C. § 103(a) rejection for claims 14 and 19.

New Claims

Applicant has added new claims 26-32 in this Amendment. New claims 26-32 are allowable based at least on their dependency upon allowable claim 1. No new matter is added claims 26-32. Furthermore, claims 26-32 recite features not disclosed by Young or Byers.

For example, claim 26 requires that a new server is added to the plurality of servers during download of data. Young and Byers are silent as to this requirement. Young contemplates that all the servers are available when the client device tests which server to download from. Byers contemplates that all the servers are available when the multicast occurs. There would no way to modify Byers to include new servers since the data is multicasted by all the servers simultaneously.

For example, claim 27 requires that at least one server of the plurality of servers is removed during download of data. Young and Byers are silent as to this requirement. Neither Young, nor Byers, contemplates a system where at least one of the server is removed during download

For example, claim 29 requires that at least two of the communication channels provide network connections to a same one of the plurality of source servers. In contrast to claim 29, as recognized by the Examiner, Young fails to provide any teaching directed to multiple network connections. In contrast to claim 29, Byers only contemplates one single point-to-point connection to each of the servers. Byers does not provide any teaching or suggestion that more than one communication channel may be connected to the same server.

For example, claim 32 requires that the highest priority data is downloaded on average before the lowest priority of data. Contrary to claim 32, in Young the lowest priority data is downloaded first. For example, as asserted by the Examiner, the initial portion of the file in Young is considered to be the lowest priority data and is downloaded first. In Young the highest priority data is downloaded subsequent to the lowest priority data. For example, as asserted by the Examiner, the subsequent portion of the file in Young is considered to be the highest priority data and is downloaded subsequent to the highest priority data. However, as made clear by the amendment to claim 32, the highest priority data is downloaded on average before the lowest priority data.

CONCLUSION

All claims in this application are in condition for allowance. Applicant respectfully requests reconsideration and prompt allowance of all pending claims. Please charge any additional fees or credit any overpayment to deposit account number 50-1778. The Examiner is invited to telephone the below-signed attorney to discuss this application.

Date: February 6, 2009

By: /Kent J. Sieffert/

SHUMAKER & SIEFFERT, P.A.
1625 Radio Drive, Suite 300
Woodbury, Minnesota 55125
Telephone: 651.286.8341
Facsimile: 651.735.1102

Name: Kent J. Sieffert
Reg. No.: 41,312